

ABSTRACT

A computer system splits a data space to partition data between processors or processes. The data space may be split into sub-regions which need not be orthogonal to the axes defined by the data space's parameters, using a decision tree. The decision tree can have neural networks in each of its non-terminal nodes that are trained on, and are used to partition, training data. Each terminal, or leaf, node can have a hidden layer neural network trained on the training data that reaches the terminal node. The training of the non-terminal nodes' neural networks can be performed on one processor and the training of the leaf nodes' neural networks can be run on separate processors. Different target values can be used for the training of the networks of different non-terminal nodes. The non-terminal node networks may be hidden layer neural networks. Each non-terminal node automatically may send a desired ratio of the training records it receives to each of its child nodes, so the leaf node networks each receives approximately the same number of training records. The system may automatically configures the tree to have a number of leaf nodes equal to the number of separate processors available to train leaf node networks. After the non-terminal and leaf node networks have been trained, the records of a large data base can be passed through the tree for classification or for estimation of certain parameter values.